

ORNL Blue Gene/P Workshop: July 29 – July 31, 2008

IBM Blue Gene/P – Optimizing, Tuning, & Scaling

Presenters:

Kirk E. Jordan, Brian Smith, Jerry Heyman, Carlos Sosa, Rajiv Bendale, Robert Walkup

IBM Cambridge, 1 Rogers Street, 3N - 3412, Cambridge., MA 02142, USA

kjordan@us.ibm.com, 1-617-693-4581

NCCS/ORNL Staff

Purpose: The purpose of this workshop is to provide an opportunity for users (application developers) to learn how to take full advantage for their codes on Blue Gene/P. This will be accomplished through some brief lectures to introduce different topics and as much hands-on effort with IBM consultants guiding the participants on their codes. This allows participants to obtain as much experience using the compilers, tools and techniques on their own codes as possible.

Day 1: The focus will be on the compilers and debuggers. Some discussion on the overall hardware and software design philosophy that is needed to understand performance issues will be given.

7:45 – 8:30 am Working Continental breakfast – sign in, setting up attendee laptops.

8:30 am – 9:00 am

☐ Introduction on how to use the local systems (Presented by ORNL staff)

○ Login and local file systems

○ Scheduling and submitting jobs

○ Policies for allocation of partitions, timing runs and scaling runs

○ I/O and file access

9:00 am – 9:45 am

☐

Blue Gene/P: Brief system and environment overview needed to understand optimization, tuning and scaling.

○ Hardware overview

☐ Systems architecture

☐ Host systems

○ Software overview

☐ Compute Node Kernel

☐ Execution process modes

☐ Message Passing Interface on Blue Gene/P

☐ Memory considerations

☐ Other considerations

Input/output

Miscellaneous

☐ Compilers overview

☐ I/O Node Software

9:45 am – 10:30 am

☐ Blue Gene/P Compiler consideration

○ Overview of the IBM compilers and linker

☐ Consideration on various compiler flags

☐ Where to start for optimization

- What options for what impact

10:30 am – 10:45 am Break

10:45 am - Noon

- ☐ Hands-on – session
- Get codes running
- Experiment with problem size
- Experiment with compiler flags/options

Noon – 1 PM Working Lunch – Question & Answer session with attendees

1:00 pm – 2:00 PM

- ☐ Debuggers overview
- Compiler and linker flags required for debugging
- GNU GDB – how to use it on Blue Gene/P
- Other debuggers at ORNL (TotalView?, need ORNL staff help here)

2:00 pm – 3:15 pm

- ☐ Hands-on - session
- Use debuggers on some simple broken code
- Experiment with debuggers on own code
- Continue problem and code development on Blue Gene/P

3:15 pm – 3:30 Break

3:30 pm – 5:00 PM

Hands-on – session 9(continued)

- Use debuggers on some simple broken code
- Experiment with debuggers on own code
- Continue problem and code development on Blue Gene/P

Day 2: The focus will be on profiling, performance monitors and on-node optimization. It is important to get on-node both core and on the 4 cores to run well. Once this is achieved scaling will then be easier.

7:45 – 8:30 am Working Continental breakfast – attendee Q&A.

8:30 am – 9:30 am

- ☐ Universal Performance Counters overview
- Hardware Description
- Low Level Software API
- Example Library and Program
- Documentation & Discussion

9:30 am – 10:30 am

- ☐ Hands-on - session
- Use UPC on some simple code
- Experiment with UPC on own kernel code

10:30 am – 10:45 am Break

Hands-on – session 9(continued)

- Use UPC on some simple code
- Experiment with UPC on own kernel code

Noon – 1 PM Working lunch – Questions & Answers with Attendees

1:00 pm – 2:00 PM

- ☐ HPCTool Kit or other Performance Tools
- Over all discussion of the HPCToolkit if available
- MPI performance: MPI Profiler/Tracer
- CPU performance: Xprofiler, HPM
- Threading performance: OpenMP profiling
- I/O performance: I/O profiling
- Visualization and analysis: PeekPerf

2:00 pm – 3:00 PM

- ☐ Hands-on - session
- Use Performance tools on some simple code
- Experiment with Performance Tools on own code
- Identify time sinks

3:00 pm – 3:15 pm Break

3:15 pm – 5:00 pm

- ☐ On-Compute Node optimization
- Information from the compiler
- Alignment for Double Hummer Performance
- Performance Discussion – On-Core, On-Node
- ☐ Profiling to identify performance issues
- ☐ Remarks on performance inhibitors
- ☐ Importance of Mapping for Torus
- ☐ Hands-on - session
- On-node optimization on some simple code
- Experiment with on own code

Day 3: The focus on scaling, tuning and libraries to improve performance

7:45 – 8:30 am Working Continental breakfast – attendee Q&A.

8:30 am – 9:30 am

- ☐ Discussion of performance and performance gains through the use of library routines
- Overview of MASS and MASSV libraries for intrinsic and math functions

- Performance improvements using ESSL routines including BLAS

9:30 am – 10:30 am

- ☐ Hands-on - session
- Use of MASS and MASSV on example code
- Use of ESSL on example code
- Experiment with on own code if appropriate, include compiler and link with libraries

10:30 am – 10:45 am Break

10:45 - Noon

- ☐ Discussion of Scaling Inhibitors
- Understanding performance and improvement through the use of good I/O techniques
- Trade offs between OpenMP and MPI on BG/P
- Other depending on the needs and interest of audience

Noon – 1 PM Working lunch - Questions & Answers with Attendees

1:00 pm – 2:45 pm

- ☐ Hands-on Session
- Continue working with own code
- ☐ Summary of Some Application Performance
- Case studies discussion on execution modes
- Some application performance
- ☐ Speculation on new approaches for applications on Blue Gene

4:00 pm end of workshop; workshop evaluation